Multi-level Modeling of Populations

Franz G. Amador & Daniel S. Weld*
Department of Computer Science and Engineering, FR-35
University of Washington
Seattle, WA 98195
franz@cs.washinton.edu & weld@cs.washinton.edu

Abstract

Many systems are too complex to represent with a single, detailed model because the model simply becomes too cumbersome. This is especially true when reasoning about populations composed of many interacting individuals. A fully-detailed model of such a system would have to model all the interactions between a potentially enormous number of individuals. Frequently, however, these interactions fall into only a few classes. We claim that when this occurs population systems are profitably modeled at three levels.

At the highest level we represent the gross properties of the population as a whole. We also provide restricted means for modeling any direct influences the population's *environment* may have upon them, such as the constraint that the volume of a gas must equal the volume of its container.

At the aggregate level, we model the statistical nature of a population explicitly. Changes to the membership of a population are modeled with sinks and sources. For example, evaporation of a liquid can be modeled as a sink with a threshold on kinetic energy of the molecules. This allows us to deduce the cooling effect of evaporation.

The most detailed level describes the properties of individual members of the population. This level includes a set of models that describe all possible *interactions* that an individual may have with other individuals and with the population's environment. For example, one model might describe a gas molecule's collisions with other molecules, while another model would represent a colision with a container wall. Changes in individual properties due to interactions at the individual level are summarized and propagated to the higher levels by statistical operators. Our approach enables powerful reasoning about aggregate properties plus the flexibility to justify conclusions in terms of individual interactions.

*This research was funded in part by National Science Foundation Grants IRI-8902010 and IRI-8957302, and a cash donation from the Xerox corporation.

1 Introduction

Modeling is a central issue in qualitative physics because the performance of reasoning algorithms is directly dependent on the size of the underlying model. For many complex systems, it is not practical to use a single monolithic model — reasoning becomes intractable and the explanations incomprehensible. In these cases it is preferable to model the system at several levels of detail, using fine-grained models only when necessary.

It is particularly important to avoid excessive detail when reasoning about populations composed of many interacting individuals. A fully-detailed model of such a system would have to model every one of the potentially enormous number of individuals and all of the N^2 interactions. Fortunately, however, these interactions can often be grouped into just a few classes. In these cases, we assert that population systems are more profitably modeled at three levels:

- Models of the classes of interaction between individuals,
- Probabilistic models of the aggregate effect of many such interactions, and
- A model of the macroscopic nature of the population as a whole.

By summarizing the results of individual level reasoning into the ontology of the aggregate and macro levels, a program can exploit the best of both models. Such multi-level reasoning can be applied to answer questions about many interesting systems. For example, we can answer the question "Why does a liquid cool as it evaporates?" by observing that evaporation involves loss of the most energetic molecules since these are the ones that can break the electrostatic bonds confining them to the liquid. Thus the average kinetic energy of the population of remaining molecules must decrease. Since

temperature is equivalent to mean molecular kinetic energy, the liquid cools down. Similarly, we can answer the question "Why do warnings about cholesterol raise the price of fish?" by noting that more individuals will want to buy fish, so the market equilibrium shifts to a higher price. Lastly, we consider the question "Why does the temperature of a gas in a cylinder increase when the gas is compressed by a piston?" It does so because molecules that collide with the inward-moving piston face are accelerated, raising their kinetic energy; this raises the average thermal energy for the population. In each of these examples we can explain large-scale effects by reasoning about only a few qualitatively-distinct individual behaviors.

Our research has resulted in a comprehensive framework for multi-level modeling of populations. At the macro level we represent the gross properties of the population as a whole. We also provide restricted means for modeling any direct influences the population's *environment* (by which we mean everything in the world besides the population) may have upon them, such as the constraint that the volume of a gas must equal the volume of its container.

In this paper, we concentrate upon the aggregate and individual levels. The aggregate level presents statistical models of a population and uses sources and sinks to represent changes to the population's membership.

The individual level includes a set of models that describe all possible *interactions* that an individual may have with other individuals and with the population's environment, such as a gas molecule's collisions with other molecules and with the container walls. Changes in individual properties due to these interactions are summarized and propagated to the higher-level models by *statistical operators* such as **Summation** and **Mean**.

Our research has focussed upon the use of multi-level reasoning to determine how population parameters change in response to a particular environment. Such reasoning allows us to predict population-level effects, and it allows us to produce more satisfying explanations of population phenomena than could be produced without reference to the individuals that make up the population. We have tested our framework with a common lisp implementation which is adequate for reasoning about systems in which the population does not gain or lose individuals, such as the gas-in-the-cylinder example above. We plan to extend the program to handle the other two examples in which the population membership changes dynamically.

2 The macro level

The macro level describes the population as a whole with its own characteristic properties, which we represent as macro-level quantities. Many macro-level quantities are defined in terms of the quantities that characterize the statistical entities of the aggregate level and are described in the next section. Some, however, are most logically defined strictly at the macro level. The volume occupied by the contained gas in our example is one such quantity. It is determined not by any aggegate behavior or state of the individual molecules, but rather by the volume of its container. We thus must be able to reason about the macro-level environment, and we accordingly allow a simple constraint model of it. The environment may have its own macro-level quantities, such as container volume, and it is composed of *components* which may be defined to influence these quantities. Thus the piston component of the contained gas system is defined to negatively influence the volume of the inside of the cylinder. We may constrain macro-level population quantities to be proportional to macro-level environmental quantities, which allows us to relate the gas volume to the cylinder volume. We have so far restricted modeling at the macro level to the bare minimum complexity in order to focus on population modeling issues. However, in the long term, we hope to investigate multi-level grounding of complex QP theory macro envisionments.

3 The aggregate level

The aggregate level allows us to relate macro-level quantities to the population of individual-level quantities. Although this population is of a finite number of individual values, it is useful to represent it in continuous terms. Accordingly, we define the *quantity-population* for a quantity Q to be a continuous function from the individual-level quantity Q to the real numbers. If f is the quantity population for Q then the number of individuals having a value between q_1 and q_2 is defined as $|\int_{q_1}^{q_2} f(q)dq|$.

Although the quantitative notion of quantity-population is powerful, we seek a qualitative description of populations. We this via a set of *statistical operators* that are defined on quantity-populations. These allow us to describe and summarize the population in high-level, qualitative terms. So far we have defined the following operators:

- **NumIndividuals**, which in discrete terms represents the size of the population. We define **NumIndividuals**(*quantity*) to be the integral from minus infinity to infinity of the quantity-population defined by the argument. Since the value for this integral is independent of the quantity concerned, we often drop the argument and refer to **NumIndividuals** as a variable rather than a function.
- **Summation**(*quantity*), which is equal to the sum of *quantity* over all the individuals in the population.
- **Mean**(*quantity*), which is equal to Summation(*quantity*) divided by NumIndividuals.
- **Min**(*quantity*), which is equal to the smallest value of the quantity which is held by an individual.

- **Max**(*quantity*), which is equal to the largest value of the quantity which is held by an individual.

Macro-level quantities may be defined to be proportional (or inversely proportional) to any of these. They may also be said to be proportional to other macro-level quantities, which allows some quantities to be "hybrid," that is, influenced both by individual behavior and by the macro-level environment. Such is the case with the density (number of molecules per unit volume) of our contained gas, which depends upon both **NumIndividuals** and volume.

Although we have implemented only a few statistical operators, they have surprisingly broad applicability. For example, we may describe the contained gas system in the following manner. [1]

```
(Qprop heat (Summation kinetic-energy))
(Qprop temperature (Mean kinetic-energy))
(Qprop density NumIndividuals)
(Qprop- density volume)
(Qprop pressure temperature)
(Qprop pressure density)
```

Here heat, a macro-level quantity, is defined to be proportional to the aggregate-level summation of kinetic energy over all individuals. Temperature is defined to be proportional to mean kinetic energy, which makes it proportional to heat and inversely proportional to the size of the population. Density and pressure are both hybrid quantities. Density is proportional to the size of the population and inversely proportional to volume. Pressure is proportional to both temperature and density. [2]

3.1 Dynamically-changing populations

Population reasoning is especially interesting when the populations can gain or lose individuals over time. Our approach for modeling this is to include special *sources* and *sinks* in the environment. A source adds individuals to a population, and a sink removes them. A source and a sink can be associated to form a *portal*, which allows individuals to move from one population to another. Thus we can represent a vessel containing both liquid and gas phases of a substance as two populations connected by two portals. One portal represents evaporation. It appears as a sink to the population of molecules confined to the liquid and as a source to the population of gas molecules. The other portal represents condensation, and appears as a sink to the gas and a source to the liquid. An alternative model might dispense with the portals; instead evaporation would be modeled with an unconnected sink on the liquid population and condensation would be represented as a simple source to the liquid. This model encodes the assumption that the mass of liquid is negligible compared to that of the gas.

The mechanism by which an individual enters or leaves a population must be modeled at the individual level, but we can still define a source or a sink strictly at the aggregate level in terms of probability distributions. For a source, the distribution represents the probability (per unit time) that an individual with any given combination of quantity values will enter the population. The distribution for a sink represents the corresponding probability of leaving. Most generally, the distribution must be multi-dimensional to assign a probability to all possible combinations of individual quantity values, but in practice, concerning ourselves with only one individual quantity generally suffices. For example, a sink representing evaporation from a liquid can be characterized by a very simple probability distribution that specifies that any individual molecule whose kinetic energy is above the breaking strength of intermolecular bonds will very likely leave the population, while those molecules with less kinetic energy will not.

We refer to a sink like this one representing evaporation as a *high-pass sink* and define it as a triple *(Q, T, R)*, where Q is an individual-level quantity, T is a threshold value for that quantity, and R is a rate of disappearance per unit time for individuals whose value for the quantity is above the threshold. We define *low-pass sinks* analogously.

3.2 Reasoning using sinks

Determining the effects of even simple sinks is not straightforward. We cannot, for example, simply define the effect of a simple sink in terms of its effects upon **NumIndividuals** or **Summation**. Since most sinks will cause both **NumIndividuals** and **Summation** to decline, the obvious derivation of the qualitative effect on **Mean** is ambiguous. This is not surprising, because what makes a high-pass sink different from a low-pass sink, say, is the differing selectivity in which individuals are removed. Reasoning about **NumIndividuals** and **Summation** separately ignores this selectivity.

Instead, we must reason about the effect of sinks directly upon the mean of a quantity. We present four rules that describe the effects of a high-pass sink upon the mean of a qualitative population:

1. If the sink threshold T is greater than **Max**(Q), there is no effect upon **Mean**(Q).

[1] Following QP theory, *Qprop* means qualitatively proportional, and *Qprop-* is inverse qualitative proportionality.

[2] We justify this definition of pressure by observing that the pressure a gas exerts upon an object is caused by the collisions of the gas molecules with the object's surface. It is thus proportional to the momentum of each colliding molecule, which varies with velocity and thus kinetic energy and thus temperature, and to the number of collisions per area per time. The number of collisions over a given area during a given time interval depends upon the number of molecules that are within striking distance of the surface at the beginning of the time interval, and this depends upon density. (It also depends upon molecule velocity, which determines how far from the surface a molecule may be and still be within striking distance during the time interval, but this too varies with kinetic energy and thus temperature.)

2. If T is less than **Max**(Q), but greater than **Min**(Q), **Mean**(Q) tends to decrease.
3. If T is less than **Min**(Q), there is no effect upon **Mean**(Q).
4. The closer T is to **Mean**(Q), the greater is the tendency for **Mean**(Q) to decrease.

The intuition behind these rules is as follows. In a population of discrete values, removing an individual value above the mean causes the mean to decrease, while removing a value equal to the mean does not affect the mean, and removing a value below the mean causes the mean to increase. Deleting a value has a greater effect upon the mean the farther it is above or below the mean.

If the threshold is above the high extreme of the population, no values are deleted, and the mean is not affected. If the threshold is above (or equal to) the mean, only values above (or equal to) the mean are removed, so the mean must decrease. The closer the threshold is to the mean the more such values are removed, so the greater the effect is upon the mean. If the threshold is below the mean but near to it, many values above the mean and a few below it are deleted; the mean still decreases, but the effect is moderated by the deletion of individuals below the mean. As the threshold moves farther below the mean, this moderating effect gets stronger, thus further weakening the decrease of the mean, until the threshold is below the entire distribution, at which point the two effects cancel, and the mean does not change.

Rule 2 above, together with definitions of evaporation as a simple, high-pass sink and of temperature as the mean of molecular kinetic energy, allows us to deduce that evaporation of a liquid will tend to cause its temperature to decrease.

We may describe low-pass sinks with an analogous set of rules, save that such sinks tend to cause the mean to increase, rather than decrease.

3.3 Reasoning using sources

It is perhaps surprising that sources are substantially different than inverse sinks. To characterize a source one must define the distribution from which the source is producing values. This is unnecessary for sinks because they bootstrap off the population distribution on which they are defined. On the other hand, it is unnecessary to specify an individual-level quantity or threshold for a source, since the population specifies these values implicitly. Thus a source can be defined as a pair (P, R) where P is a population and R is a rate of production: individuals per unit time.

A *high-pass portal* (a linked high-pass sink / source pair) is defined the same as a high-pass sink, as a triple (Q, T, R) where Q is an individual-level quantity, T is a threshold value for that quantity, and R is the rate of transfer. In this case, the source's population is derived from that of the sink relative to the specified quantity and threshold.

4 The individual level

The individual (or micro) modeling level describes the individuals that make up the population and their interactions with the environment of the population and with one another. The state of an individual is described by a set of quantities. In the gas-in-the-cylinder example, we describe the state of an individual molecule by its position, velocity, and kinetic energy. The value of a quantity may be affected by an individual's interactions with other individuals and by the individual's interactions with the environment, such as the cylinder wall. To be able to reason about changes in the values of the population's properties we must model all such interactions.

For any given system, we may group these interactions into *types*. For example, each object with which a molecule can collide requires a different type of collision interaction. For simplicity and tractability we model and reason about each type of interaction separately, rather than attempting to model all possible sequences of interactions that an individual may engage in. This allows us to factor the problem of micro-level modeling into simpler sub-problems. Doing so requires that we assume that interactions are independent, i.e. the occurrence of one will not influence the course of another. This assumption is justified so long as the course of one interaction does not depend upon unmodeled "hidden state" that a participating individual carries from prior interactions, or, in other words, that our representation of an individual captures all relevant aspects of its state. The validity of our assumption also depends upon the occurrence of one interaction not interfering with another while it is in progress, but this is simply a closed-world assumption: we assume that the model of each interaction fully describes all possible outcomes.

Thus we require a separate model for each possible interaction type. At least one of these models must describe the interaction of an individual with another individual. For the contained gas example, this means the collision of two molecules. Other models describe the interactions of individuals with the environment. We assume that the overall (*macro*) environment of the population is large relative to an individual, so that the individual interacts with only a small part of it, which we term the *micro-environment* for that interaction. Different parts of the macro-environment may each present a different micro-environment to the population's individuals. We define each micro-environment with a single model that describes all the ways in which an individual may interact with that micro-environment. For example, the gas-in-the-cylinder example presents two micro-environments (as well as an individual/individual interaction which is independent of the environment) as shown in figure 1:

- A fixed surface, which represents an unmoving, insulated container wall.
- An inward-moving surface, which represents the

piston's compressing action.

If the piston could also move outwards, and hot and cold surfaces were modeled as well as insulated ones, there would be a total of nine possible micro-environments.

Since we are modeling effects upon individuals entering and leaving interactions, it is natural to use a process-centered modeling ontology because it allows for a varying topology of interaction. Accordingly, we use Qualitative Process (QP) theory [Forbus, 1984] to specify the micro-level models. An individual is represented as an *entity* with *quantities*. *Processes* and *views* allow direct and indirect *influences* upon its quantities. For example, a motion process allows velocity to influence position, and force processes allow forces to create accelerations, which are summed to influence velocity.

An interaction that involves more than one individual may act to conserve some individual-level quantity, which means that its sum over all the participating individuals cannot change. In the contained gas example, the molecule-molecule collision interaction conserves the sum of the kinetic energies of the two individuals. As will be discussed further below, knowing about such conserved quantities can be quite useful, so in addition to their QP theory description, interaction models may have a list of conserved quantities.

4.1 Individual-level modeling of sources and sinks

While the aggregate-level description of sources and sinks allows reasoning about their effects upon the population, more satisfying explanations of their behavior can be generated if we have models of the mechanisms underlying them. Such models logically reside at the individual level. For example, we could associate with a sink that represents evaporation a model that represents a molecule's attachments to its neighbors via spring-like bonds that will break if the molecule vibrates too energetically.

5 Aggregating individual behavior

Recall that our goal is to be able to determine whether and in what direction macro-level quantities change given a particular environment and to be able to explain our reasoning. This is equivalent to asking for the signs of their derivatives. How we go about calculating these signs depends upon how each quantity is defined.

5.1 Quantities influenced by the macro-environment

The calculation is easy if the quantity is defined solely in terms of influences from macro-environment components — a qualitative sum of of the influences determines the qualitative derivative. For example, we may determine the derivative of volume this way for the contained gas system. Since gas volume is proportional to cylinder volume, and the only influence upon that is the negative influence from the piston component, we can conclude that gas volume is decreasing.

Doing so requires assuming that there are no other, unstated influences (a closed-world assumption). Note that since we are summing qualitative values, it is possible to get an ambiguous answer, which simply means that it is consistent with the available information for the quantity to be increasing, decreasing, or remaining constant.

5.2 Quantities influenced by the micro-environment

Those quantities that are defined in terms of the statistical operators require more elaborate computation. First we consider **NumIndividuals**. Obviously, sources tend to cause the population size to increase, and sinks tend to cause it to decrease. If both are present, the effect will be ambiguous, although we may be able to determine the existence an equilibrium state in which the population size remains constant. If there are no sources or sinks, the derivative for **NumIndividuals** is zero, as is the case with our contained gas example system.

To determine the derivative of a macro-level quantity defined using **Summation** with perfect accuracy would require determining the derivative signs of the corresponding individual-level quantity for every individual in the population and summing them. For populations with a large number of individuals (like a gas) this is an unworkable idea, so we must make some assumptions that allow us to reduce the problem to a manageable size. We assume that the set of micro-level models appropriate to the environment represents all possible interactions that can influence individual-level quantities (another closed-world assumption). We also assume that each such interaction will actually occur for at least some individuals. This assumption is based upon the idea that there are so many individuals in the population that some will at one time or another be in every possible state and thus will be capable of entering into every possible interaction.

Given these assumptions, we calculate the derivative sign for **Summation** in in three steps: 1) generate a total envisionment for each micro-environment and isolate all possible behavior paths, 2) compute the *net change* for the quantity over each behavior, and 3) sum the net changes together.

5.2.1 From envisionments to behaviors

A total envisionment [Forbus, 1984] is a directed graph composed of all possible qualitative states of a system along with all legal transitions between them. Each state includes the qualitative values and qualitative derivatives for every quantity. Following [Forbus, 1987], we call a state with no incoming transitions an *eden state*. We also define a *dead-end state* to be one with no outgoing transitions and a *behavior* to be a path from an eden state to a dead-end state. We assume that every such behavior is physically realizable.

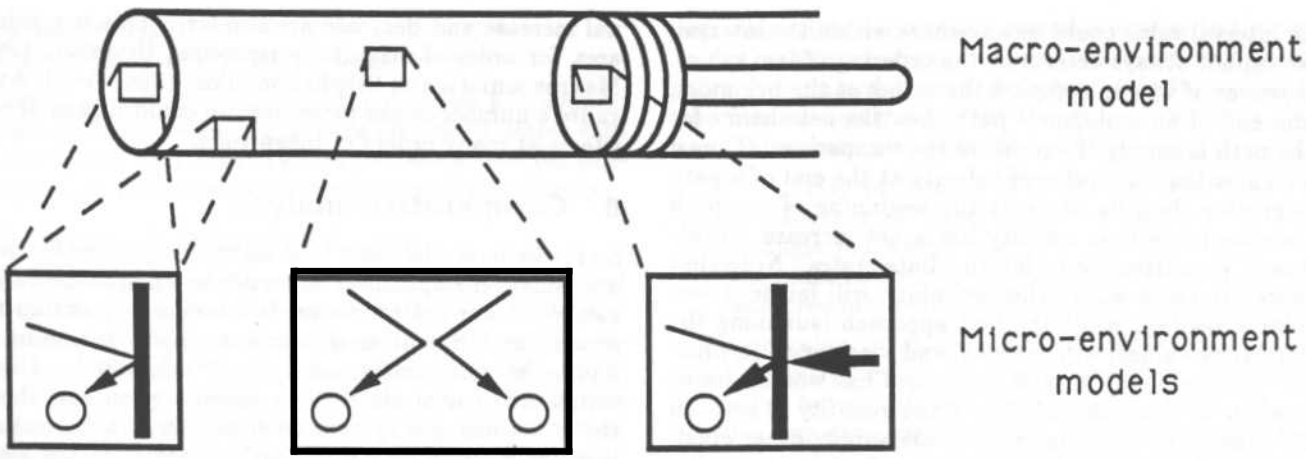

Figure 1: Relation of micro and macro environments. The macro-environment is divided into components. Each micro-environment model describes the interaction of an individual with a micro-environment. Components and micro-environments need not correspond in a one-to-one manner.

This amounts to assuming that our qualitative state description is sufficiently discriminating to prevent spurious behaviors.

The intuition here is that the set of behaviors represents all the possible things that may happen to an individual that participates in the modeled interaction. We are thus assuming that the individual must enter the interaction at an eden state and exit at a dead-end state. This is based upon a deeper assumption that all interactions begin and end within a finite interval of time, that is, we do not allow interactions that continuously influence an individual. [3] Thus, barring cyclic envisionments, there must be at least one state that represents the individual before it enters into the interaction and one that represents the individual after it leaves the interaction, and these must be eden and dead-end states, respectively.

Note that an envisionment with one or more cycles may have no eden or dead-end states and hence no behaviors by our definition. Furthermore, even if there are eden and dead-end states, the presence of a cycle means there may be an infinite number of paths between some of them. We note the following:

- A cycle in which all the states have the same derivative value for the quantity can be collapsed into a single state with that derivative. This can create new eden or dead-end states.
- When analyzing behaviors in a cyclic envisionment, going around a cycle twice or more cannot yield a new net change value for any parameter. Thus there are only a finite number of "significant" behaviors from the perspective of determining net change values.

[3] One can easily imagine such interactions, for example, the influence of the magnetic field upon particles in a ring-shaped accelerator, but currently we leave such systems as yet another topic of future research.

5.2.2 Computing the net change for a behavior

For each significant behavior in each envisionment, we need to determine the net effect on the quantity of performing the behavior. We have developed three techniques for computing the net effect: summing the quantity's derivative, comparing landmark values, and using a reference quantity. In order to minimize ambiguity in some examples, all three techniques must be used.

In the simplest case, we can determine this net change by summing the signs of its derivative in each state along a behavior. This analysis is all that is necessary for our gas-in-the-cylinder example. Each model of the collision of a molecule with a surface yields a single behavior in which there are states for the molecule's approach, contact with the surface, a spring-like compression and decompression, and the molecule's flying away. While the molecule is in contact with the surface its energy (defined as the sum of kinetic energy and compressional potential energy) changes monotonically. Thus it is easy to determine that collision with a fixed surface causes kinetic energy to remain unchanged, that collision with an inward-moving or hot surface causes kinetic energy to increase, and that collision with an outward-moving or cold surface causes kinetic energy to decrease.

However, if we find both a positive and a negative derivative along a behavior, the sum will be ambiguous. One way to resolve the ambiguity is to use information from the quantity's quantity space [Forbus, 1984], which is a partially-ordered set of landmark points. At any time the quantity's value is either equal to a landmark point or in the interval between two landmarks. Because the landmarks are only partially ordered, and because

an interval value could be anywhere within the interval, we cannot always determine the ordering of two values. However, if we can compare the values at the beginning and end of an ambiguous path then the net change for the path is simply the result of the comparison. Thus if we know that the value of velocity at the end of a path is greater than its value at the beginning of the path then we know that velocity has a net increase regardless of its derivative in intermediate states. Note that there are cases where this technique will fail to determine a relation while the first approach (summing the derivative values) will succeed, and vice versa. In practice, the only way in which this idea has worked for us is when we know the relation of the quantity to zero. In QP theory, all other landmarks are simply other quantities, and we cannot know *a priori* whether those other quantities are constant.

Even if we cannot order the values in the beginning and ending states, we can still resolve ambiguity in some cases. If a quantity is known to be equal to another quantity at the beginning and end of a path, then the net change values for the two quantities must be the same, so if either is known then both are known. Also, if the quantity is less than the reference quantity at the beginning and greater than (or equal to) it at the end (or equal to it at the the beginning and greater than it at the end), and the reference quantity has remained constant or increased, then we can conclude that the quantity of interest has increased. Similarly, if the quantity has decreased relative to the reference quantity, and the reference quantity has remained constant or decreased, then we can conclude that the quantity of interest has decreased.

5.2.3 Summing net changes

To return a value for the **Summation** expression we must add the net change values for every behavior in every envisionment. Obviously the potential for ambiguity is large if there are many possible behaviors. One way to reduce this potential is to know which interactions conserve the quantity being summed. Because we know that they cannot affect the quantity sum for the individuals that participate in them, we can ignore them. This fact is critical to the successful analysis of systems such as the gas-in-the-cylinder example. When two molecules collide, the kinetic energy of each may increase or decrease. Thus if we did not know that this interaction conserves kinetic energy, any summation of kinetic energy would always be ambiguous.

What is really needed is some means of reasoning about the probability of each possible individual behavior. This would allow us to determine the relative number of individuals whose quantity has each possible change value. If we know that there are significantly more individuals for which the quantity increases than there are for which it decreases then we can unambiguously say that the summation increases (assuming that the amounts of each individual increase and decrease are similar). This is a prime area for order-of-magnitude reasoning [Raiman, 1986, Mavrovouniotis and Stephanopoulos, 1988]. With Avogadro's number of gas molecules we could reason about effects at many orders of magnitude.

6 Comparative analysis

So far we have discussed how we can determine the way a population responds to a particular environment. We also want our system to predict how a population responds to different environments. Such reasoning is known as comparative analysis [Weld, 1990]. Using comparative analysis we can answer questions about the contained gas system such as "Does temperature increase more rapidly if the piston moves inward more rapidly?"

For now we consider only comparisons between environments whose differences can be expressed as differential changes in the value of one or more quantities. These quantities can be either part of the macro model (as with the question about piston speed above) or about the nature of the individuals (e.g., "What would happen to pressure if the molecular weight of the gas were higher?").

The overall algorithm requires three phases. First, if the comparative analysis question specifies a change to an aggregate or macro-level quantity, this is converted into the corresponding change at the micro level. Second, comparative analysis is performed in the individual models. Third, the relative change (RC) results predicted by comparative analysis are aggregated to refer to the quantities of the initial model. Since this aggregation process is performed using the statistical operators **NumIndividuals**, **Summation**, and **Mean** in the same way that ordinary derivative values are summarized, we concentrate here on how comparative analysis can be performed on individual-level models. We have discovered that generalizing DQ analysis [Weld, 1988] to work on QPE envisionments instead of QSIM behaviors [Kuipers, 1986], requires three extensions (in addition to the process of extracting behaviors from envisonments described in section 5.2.1). These extensions are described in the sections below:

6.1 Constraint Compilation

For each state in the behavior, the process and view structure must be compiled into the corresponding qualitative differential equation (e.g., QSIM constraints); these constraints are then asserted into an AMORD-like database [de Kleer *et al.*, 1977] as input for the DQ analysis inference rules. A closed-world assumption is used to sum influences from the active processes to determine quantity derivatives. For example if a state had two force processes active on a single individual `X`, each influencing its acceleration, the following QSIM constraints would be generated:

```
(ADD X-acceleration-001 X-acceleration-002
```

X-Deriv-059)
(D/DT X-velocity X-Deriv-059)

where X-acceleration-001 denotes the acceleration quantity from the forcing process instance 1 active with individual X, and X-Deriv-059 denotes a new parameter created to represent the derivative of X's velocity.

6.2 Multiple Operating Regions

The existing DQ analysis rules assume that a single set of qualitative differential equations holds for all time, but this is not the case for behaviors with changing process and view structures such as result from collisions in the contained gas example. Fortunately, it is a simple matter to add an extra argument in the rules that identifies the appropriate QDE constraints as a function of each state's process and view structure.

6.3 Identify Exogenous Parameters

To perform complete comparative analysis it is necessary to know which parameters are exogenous so that an initial relative change value may be ascribed to them. When performing comparative analysis on a QSIM model, this information must be given by the human modeler. However, since QP theory embodies a theory of causality [Forbus and Gentner, 1986], much of this information can be deduced from the QP models. The remainder can be deduced from the links between models. For example, if the macro-level question specifies that the piston is moving faster then the initial relative change for the micro-model must specify that the wall speed is moving faster yet the wall heat is unchanged.

7 Related work

This work build upon and extends that of many other researchers. [Collins and Gentner, 1983] describes psychological studies of how subjects reasoned about evaporation using multiple models much like we have described here. Perhaps the best way to view this work is as an elaboration and implementation of their initial studies.

Our previous work on "aggregation" [Weld, 1986] is similar to this research, but there are several differences. While our current approach uses a fixed hierarchy of three levels, aggregation created abstract models of system behavior dynamically and was not limited to three levels. However, aggregation could only summarize the effects of repeating cycles of processes. Furthermore, aggregation could not summarize the effects of random processes — our work here uses more sophisticated statistical representations.

Recent work on the "molecular collection" ontology [Collins and Forbus, 1987] is also relevant to our endeavor. Although Hayes first made the case that a "contained substance" ontology was insufficient by itself for commonsense reasoning [Hayes, 1985], Collins and Forbus implemented a qualitative simulator that used both representations. Our macro model is similar to the contained substance view, but our molecular models are more fine-grained than a molecular colection. Collins and Forbus neatly observed the distinction when they say that a molecular collection is a

> ... tiny piece of stuff [which] is viewed as a collection of molecules — as opposed to a single molecule — so that it may possess such macroscopic properties as temperature and pressure. [Collins and Forbus, 1987, p590]

By choosing to model a population in terms of its most primitive members, the difficulty of integrating reasoning between the levels is greatly increased, motivating our development of the statistical aggregation operators which are unnecessary with the molecular collection ontology. We feel that both approaches are useful for reasoning, especially about complex systems such as refrigerators and heat engines. For example, a molecular collection could be used to select a sequence of micro-environments for individual-level analysis. We hope to unify these approaches as part of our future work.

With motivation very similar to ours, Liu and Farley observe that the device ontology is incapable of answering certain questions about electronics such as "Why does the current through a resistor increase when the voltage across it increases?" To rectify this explanatory inadequacy, they introduce the "charge-carrier" (CC) ontology [Liu, 1989, Liu and Farley, 1990]. The CC ontology models electronics with primitives such as "field" and "charge flow" rather than "resistance" which is modeled in the device ontology. Liu and Farley use "bridging relations" to link the two ontologies, but unlike our approach, these links are not statistical in nature.

8 Conclusion

We have presented a novel way of reasoning about systems that can be viewed as populations of individuals. Our framework models the population and its environment at the levels of the macroscopic system as a whole, the aggregate properties of the population of individuals, and at the level of the individual themselves. At the micro level we model all possible interactions between an individual and other individuals and between an individual and its environment. We determine the effects of these interactions upon the individuals by analyzing the envisionment generated from each model. We summarize these effects to the aggregate level by means of statistical operators. This allows us to determine the derivatives of population quantities under the influence of an environment. At the aggregate level we also consider the influence of sources and sinks which allow the population to gain and lose members. All these conclusions are propagated to the macro-level quantities through qualitative proportionalities. Adding comparative analysis to our system will allow us to increase

the range of questions that can be answered to include hypothetical changes to the environment.

8.1 Implementation

We have implemented a subset of the framework we describe here. This includes representations for macro, aggregate, and micro level models, rules that implement the statistical operators and macro-environment influences and qualitative proportionalities, and a program that implements the above-discussed algorithms for determining net change values from an acyclic envisionment. Our program is written in COMMON LISP and we use Forbus' QPE [Forbus, 1989] to generate envisionments from the micro-environment models. The rule-based reasoning is done by RULER, a simple inference engine with an associated TMS based on AMORD [de Kleer *et al.*, 1977]. The gas-in-the-cylinder example is implemented and works correctly. We are just starting to implement comparative analysis for QPE envisionments; given the availability of a compiler from QP theory models to QSIM qualitative differential equations [Crawford *et al.*, 1990], we may switch our efforts in this direction.

8.2 Future Work

Our current implementation can reason only from micro to macro level, following our intuitions about causality. However, the proportionalities and statistical operators are probably best implemented as constraints, which would allow reasoning in the reverse direction as well. We could then produce answers to a question such as "What can cause the temperature to increase?"

A good theory of sources and sinks should also allow us to predict various equilibria. For example, we should be able to deduce that a liquid/gas system in a closed container will have equilibrium population sizes and that a liquid in an open container will have an equilibrium boiling point.

We are particularly interested in the behavior of some populations that continuously renormalize their quantity value distributions. For example, molecule velocities in an ideal gas always form a Maxwell-Boltzmann distribution, even if molecules are added or removed. Such self-normalization could be modeled at both the level of the inter-individual interactions that accomplish it and at the aggregate level, where we are exploring the intriguing possibility of representing self-normalization as a portal from the population to itself.

Other interesting topics include allowing the population to affect its environment, introducing new statistical operators (a good candidate is **Variance**), reasoning about several interacting populations, reasoning about subpopulations within populations, and introducing behavior probabilities and order-of-magnitude measures of the number of individuals affected in each possible way.

ACKNOWLEDGEMENTS

This work has benefited from conversations with Steve Hanks, Dorothy Neville, and Scott Penberthy. This research was funded in part by National Science Foundation Grants IRI-8902010 and IRI-8957302, and a donation from the Xerox corporation.

References

[Collins and Forbus, 1987] J. Collins and K. Forbus. Reasoning About Fluids Via Molecular Collections. In *Proceedings of AAAI-87*, July 1987.

[Collins and Gentner, 1983] A. Collins and D. Gentner. Multiple Models of Evaporation Processes. In *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*. Erlbaum, May 1983.

[Crawford *et al.*, 1990] J. Crawford, A. Farquhar, and B. Kuipers. QPC: A Compiler from Physical Models into Qualitative Differential Equations. In *Proceedings of AAAI-90*, August 1990.

[de Kleer *et al.*, 1977] J. de Kleer, J. Doyle, G. L. Steele, and G. J. Sussman. AMORD: Explicit Control of Reasoning. In *Proceedings of the Symposium on Artificial Intelligence and Programming Languages*, Rochester, N.Y., August 1977. SIGART.

[Forbus and Gentner, 1986] K. Forbus and D. Gentner. Causal reasoning about quantities. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, pages 196–207, 1986.

[Forbus, 1984] K. Forbus. Qualitative Process Theory. *Artificial Intelligence*, 24, December 1984.

[Forbus, 1987] K. Forbus. The Logic of Occurrence. In *Proceedings of IJCAI-87*, pages 409–415, August 1987.

[Forbus, 1989] K. Forbus. The Qualitative Process Engine. In *Readings in Qualitative Reasoning about Physical Systems*, pages 220–235. Morgan Kaufmann, August 1989.

[Hayes, 1985] P. Hayes. Naive Physics I: Ontology for Liquids. In *Formal Theories of the Commonsense World*. Ablex, 1985.

[Kuipers, 1986] B. Kuipers. Qualitative Simulation. *Artificial Intelligence*, 29, September 1986.

[Liu and Farley, 1990] Z. Liu and A. Farley. Shifting Ontological Perspectives in Reasoning about Physical Systems . In *Proceedings of AAAI-90*, August 1990.

[Liu, 1989] Z. Liu. A Charge-Carrier Ontology for Reasoning about Electronics. In *Third Workshop on Qualitative Physics*, August 1989.

[Mavrovouniotis and Stephanopoulos, 1988] M. Mavrovouniotis and G. Stephanopoulos. Formal Order-of-Magnitude Reasoning in Process Engineering. *Computer Chemical Engineering*, 12:867–880, 1988.

[Raiman, 1986] O. Raiman. Order of Magnitude Reasoning. In *Proceedings of AAAI-86*, August 1986.

[Weld, 1986] D. Weld. The Use of Aggregation in Causal Simulation. *Artificial Intelligence*, 30(1), October 1986.

[Weld, 1988] D. Weld. Comparative Analysis. *Artificial Intelligence*, 36(3), October 1988.

[Weld, 1990] D. Weld. *Theories of Comparative Analysis*. MIT Press, Cambridge, MA, May 1990.